

②

DTIC FILE COPY

AD-A225 619

**KNOWLEDGE-BASED ENVIRONMENTS
FOR TEACHING AND LEARNING**

Technical Report AIP - 129

Beverly Woolf, Elliot Soloway, William Clancey,
Kurt VanLehn¹ & Dan Suthers

**The Artificial Intelligence
and Psychology Project**

Departments of
Computer Science and Psychology
Carnegie Mellon University

Learning Research and Development Center
University of Pittsburgh

DTIC
ELECTE
AUG 24 1990
S B D
Co

2

KNOWLEDGE-BASED ENVIRONMENTS FOR TEACHING AND LEARNING

Technical Report AIP - 129

Beverly Woolf, Elliot Soloway, William Clancey,
Kurt VanLehn¹ & Dan Suthers

¹Departments of Psychology
and Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213 U.S.A.

May 1990

To appear in *AI Magazine*.



¹This research was supported by the Information Sciences Division, Office of Naval Research, under contract number N00014-86-K-0678 and by the Personnel and Training Research Program, Psychological Sciences Division, Office of Naval Research, under contract N00014-88-K-0086. Reproduction in whole or in part is permitted for any purpose of the United States Government. Approved for public release; distribution unlimited.

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; Distribution unlimited		
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE					
4. PERFORMING ORGANIZATION REPORT NUMBER(S) AIP - 129			5. MONITORING ORGANIZATION REPORT NUMBER(S)		
6a. NAME OF PERFORMING ORGANIZATION Carnegie Mellon University		6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Computer Sciences Division Office of Naval Research (Code 1133)		
6c. ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, PA 15213			7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, VA 22217-5000		
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Same as Monitoring Organization		8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0678		
8c. ADDRESS (City, State, and ZIP Code)			10. SOURCE OF FUNDING NUMBERS p40005ub201/7-4-86		
			PROGRAM ELEMENT NO N/A	PROJECT NO N/A	TASK NO N/A
			WORK UNIT ACCESSION NO N/A		
11. TITLE (Include Security Classification) Knowledge-based environments for teaching and learning					
12. PERSONAL AUTHOR(S) Beverly Woolf, Elliot Soloway, William Clancy, Kurt VanLehn, & Dan Suthers					
13a. TYPE OF REPORT Technical		13b. TIME COVERED FROM 86Sept15 TO 91Sept14		14. DATE OF REPORT (Year, Month, Day) 1990July3	
15. PAGE COUNT 7					
16. SUPPLEMENTARY NOTATION To appear in the Fall 1990 issue of <u>AI Magazine</u> , published by the American Association for Artificial Intelligence, Palo Alto, CA.					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP	cognitive modelling, human-computer interaction		
			learning		
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
<p>✓ The Spring Symposium on Knowledge-based Environments for Teaching and Learning focused on the use of technology to facilitate learning, training, teaching, counseling, coaxing and coaching. Sixty participants from academia and industry assessed progress made to date and speculated on new tools for building second generation systems.</p> <p>Selection of topics and participants was motivated by a desire for ideological breadth and depth. Panel leaders included William J. Clancey and Alan Lesgold (researchers of real-world systems); Kurt VanLehn (champion of cognitive models); Beverly Woolf (defender of discourse systems); Elliot Soloway (advocate for alternative environments); and Sarah Douglas (spokesperson for supportive systems).</p>					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION		
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Alan L. Meyrowitz			22b. TELEPHONE (Include Area Code) (202) 696-4302		22c. OFFICE SYMBOL N00014

To appear in the Fall 1990 issue of AI Magazine, published by the American Association for Artificial Intelligence, Palo Alto, CA.

Knowledge-based Environments for Teaching and Learning¹

Beverly Woolf
Elliot Soloway
William J. Clancey
Kurt VanLehn²
Dan Suthers

July 3, 1990

Contact Person:
Beverly Woolf, Computer and Information Science,
University of Massachusetts, Amherst, Mass. 01003
Tel: 413 549 6036, EMail: Bev@cs.umass.edu

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

Second Generation Systems

The Spring Symposium on Knowledge-based Environments for Teaching and Learning focused on the use of technology to facilitate learning, training, teaching, counseling, coaxing and coaching. Sixty participants from academia and industry assessed progress made to date and speculated on new tools for building second generation systems.

Selection of topics and participants was motivated by a desire for ideological breadth and depth. Panel leaders included William J. Clancey and Alan Lesgold (researchers of real-world systems); Kurt VanLehn (champion of cognitive models); Beverly Woolf (defender of discourse systems); Elliot Soloway (advocate for alternative environments); and Sarah Douglas (spokesperson for supportive systems).

Human-Computer Interaction

Researchers have moved away from building omniscient tutors capable of detecting all possible errors and misconceptions. Instead, research is now focused on building empathetic partners that choose from among several forms of interaction based on the content of the communication and the needs of the student [Woolf, 1988]. Possible communication styles include didactic

¹Beverly Woolf and Dan Suthers have received partial support from the National Science Foundation under grant number MDR 8751362, and from External Research, Apple Computer, Inc., Cupertino, CA.

²Kurt VanLehn's research is supported by the Office of Naval Research under contracts N00014-88-K-0006 and N00014-86-K-0676.

explanation, guided discovery learning, coaching or coaxing, and critiquing. Although no one style is preferred, different tutorial applications will be better addressed with a given primary style.

For example, as explained by Dan Suthers and James Lester, didactic explanation is good for communicating a body of declarative knowledge shared by some community (e.g. biologists). In such applications, the student needs to learn the community's terminology, and thus didactic explanation may be more efficient than requiring that the student rediscover the principles of the field on his or her own. On the other hand, the more active nature of discovery learning helps the student "own" the acquired knowledge to a greater extent than can didactic explanation.

The style of interaction varies *within* a tutorial domain as well as across types of domains. For example, Lewis et al. [1990] showed how a (human) guided discovery tutor changed strategies from script-like to opportunistic when students suggested an activity or showed the need for remediation of a deficiency.

Communication Research Issues

Pressing research issues in human-computer communication were identified both in *artificial intelligence* and in *education*. In *artificial intelligence*, research issues include the representation and control of knowledge. From this perspective, knowledge of didactic explanation might be represented and organized in a system, along with the basic knowledge of a domain. Indexing mechanisms for accessing different perspectives on the topic should be designed using abstractions appropriate for the content selection task.

Choosing and organizing domain knowledge provides the next set of research issues. Control should account for the tutor's ability to dynamically switch strategies according to multiple constraints in a manner be sensitive to features that human tutors use in tutorial interactions. The tutor should consider available student modeling/diagnosis when making tutorial decisions based on multiple goals. Further work is required in characterizing "relevance" for selecting knowledge for didactic explanation, especially when multiple perspectives on the topic are available. Even when the primary emphasis is on stimulating the student's own creativity and intelligence, the program's design must still be based on solid theory of relevance to select its actions and response. To do so, memory and pragmatic knowledge should be brought to bear on language processing.

Another research issue concerns the characterization of coherence in machine response. Is coherence a property of the "knowledge pool" to be used in generating the next response or a property of the dialogue or both? In choosing content from a multiple granularity knowledge base, how do we ensure that the chosen pool of knowledge is coherent given the dialogue context?

Educational research issues focus on adequately modeling the student and the pedagogical context (see next section), and then recognizing how a system might stimulate and facilitate the student's own abilities and creativity.

A separate issue concerns how relevant knowledge should be presented once it has been selected. For example, the tutor might state generalizations, use case examples, or provide analogies. Presentations, whether explanations or examples, must be presented in such a way that the student will be prepared to understand new material and integrate it into an existing conceptual framework, or into one which has been built up in the preceding dialogue. We need to better understand how to choose and coordinate multi-media/modality presentations at the interface media level, e.g. the use of text, diagrams, charts, pictures, animation, and sound.

In summary, despite much work attempting to do so, we still need to figure out how to make dialogue sensitive to dialogue context and to what is known or knowable about the student's "state".

Cognitive Modeling

The cognitive modeling group provided strong advocacy for cognitive modeling at the symposium. They argued for increased use of modeling at three stages of design of knowledge-based systems, primarily (1) development of pedagogical and subject-matter theories, (2) design of instruction, and (3) delivery of instruction. Of these phases, the design of instruction is the one that seems to have achieved the most direct benefit from cognitive modeling, including substantial benefits from modeling subject matter experts. For instance, Anderson et al. [1990] attribute much of the success of their tutors to the cognitive task analysis of experts in Lisp, geometry and algebra.

Work on modeling good teachers and tutors has only just begun (with the exception of a few early classics, such as the work of Stevens and Collins on Socratic tutoring [1977]). VanLehn expects this line of investigation to pay off at least as well, if not better, than the modeling of experts and learners.

Of the three phases of pedagogical work, the actual delivery of the instruction is the area where cognitive modeling has found the least fruitful application. Mostly, this is due to a historical accident. In most systems to date, teacher models have been weaker than expert models and student models. Although a good teacher model might compensate for an impoverished expert or student model, experience has shown that strong expert and student models require a decent teacher model for the system to be effective.

VanLehn underscores the fact that modeling is just good engineering practice, regardless of whether one is building a hydroelectric dam or a science course. With tongue in cheek, he suggests that if students could sue malfeasant instructional developers, cognitive modeling would be much more common since it is so obviously effective.

William J. Clancey, however, was more reserved about the utility of cognitive modeling. While acknowledging that building such models is possible, he questions the relation they have to mechanisms of human learning. For instance, does the model show the student how to interpret and generate domain concepts, or does it simply justify the machine's presentations? Clancey

would like to see alternative cognitive models available within a system rather than a single 'correct' model used to justify instruction.

Understanding Plans and Goals

In the move away from building all-knowing and all-powerful tutors, researchers have focused on developing environments that implicitly elicit information about student goals and plans. Human dialogue succeeds despite ambiguity and digressions because both participants model the discourse, the subject matter, and the other speaker; and both participants actively work towards success of the discourse.

This suggests that continuing efforts be made to enhance the machine's ability to do its part. Techniques such as plan recognition and learning still play only a small role in current teaching systems. Interfaces were described that inquire about beliefs and high-level thoughts while supporting meta-cognitive activities. Students might choose from a menu of high-level plans, such as a menu item in an Algebra tutor that says "collect all variables to one side of the equation." Such interfaces require more careful analysis and structuring of the task domain and of cognitive structures; they also require mechanisms to support co-operative dialogue and to 'understand' student perspectives.

Real-World Applications

William J. Clancey and Alan Lesgold led several discussions on the impact of knowledge-based systems in industry and the military. The clear emergence of new architectures and positive training results have produced the feeling that progress is being made. Indeed, several systems were described which achieve the two-sigma effect [Bloom, 1984], which is the same improvement in learning that results from one-on-one human tutoring over classroom tutoring. Several success stories were described in which students using tutors learned knowledge and skills in one-third to one-half the time it took for a control group to learn the same material [Shute, 1990].

In one special case, students working with an Air Force electronics troubleshooting tutor for only 20 hours gained a proficiency equivalent to that of trainees with 40 months (almost 4 years) on-the-job training [Lesgold, Lajoie, Bunzo & Eggen, 1990]. In another example, students using a LISP tutor at Carnegie-Mellon University [Anderson, 1990] completed programming exercises in 30% time than those receiving traditional classroom instruction and scored 43% using a microworld environment learned general scientific inquiry skills and principles of basic economics in one-half the time required by students in a classroom setting [Shute, Glaser & Raghavan, 1989].

Given these results, the group asked why more tutors were not being used and why existing systems were not more effective. One reason why industry and the military have not widely

adopted these systems relates to the lack of artificial intelligence development tools, such as shells and frameworks, similar to the shells used to build rapidly expert systems. Tools would facilitate large-scale development; and a simple tool, such as a simulation tied to an expert system or to a lock-step tutor, might be a practical way for a designer to get started on a path of incremental design through feedback from the user. Some researchers suggested that a teacher should interact with a variety of tools, much as a director might orchestrate a suite of tools.

Other reasons for the slow adoption of new systems might include the need to reduce cognitive task analysis to engineering practice and the need to make widely available knowledge representations (i.e. qualitative simulations) which are better than those offered by first-generation expert system tools. An additional barrier is the lengthy development cycle required before a system can move from research lab to a saleable product.

'Hot' Research Issues

Several areas emerged as 'hot' or new research areas. These were discussed throughout the symposium.

Situated learning (and teaching/acting/planning) arose frequently as a topic. It was espoused primarily by William J. Clancey, Jeremy Roschelle, and Etienne Wenger all from the Institute for Research on Learning, Palo Alto. Since situations or contexts in which a skill is learned can not be exhaustively or completely described, training systems inevitably predetermine what is relevant. Similarly, conventional Artificial Intelligence models of expertise leave out how experts know what is relevant and how they change their minds. This approach suggests that Artificial Intelligence systems need to place increasing emphasis on representation as an activity within a perceptual space and organized by social interaction. Current systems omit the social context in which domain representations are created, justified, and changed. At present, knowledge-based cognitive modeling cannot characterize the work somebody must do to understand specific artifacts or tools of a community. One reason why on-the-job training is more cost efficient is because there is no need to simulate what can't be made fully explicit anyway.

Computer as mediator - Jeremy Roschelle demonstrated that a system could facilitate discussion amongst several students and could support their own explanations to each other. In such a case, the computer becomes a mediator empowering both students and teachers.

Andrea diSessa showed that the goals/capabilities to be taught are negotiable; he enables students to discover their own interests and assists them in pursuing same. For example, he showed a Boxer system which supports young students in inventing representations using graphic tools. In other words, they discovered the rules of graph construction.

Empowering curriculum designers - Jim Spohrer described systems developed at Apple to assist curriculum designers in the use of multi-media. Oliver Selfridge challenged the group to question the nature of the learning task implicit in their teaching machines.

Qualitative reasoning - Ken Forbus demonstrated that a system could qualitatively model a complex domain, e.g. a steam boiler or a propulsion plant, and that this representation could be used for teaching. His work on qualitative modeling is now 10 years old and is approaching the point of formalizing the reasoning needed in qualitative modeling.

Conclusions

Given the diverse backgrounds and methodology of participants, little commonality could have been expected. However, a small consensus was achieved and some new scientific ground broken. Agreement was reached on the need for a variety of discourse approaches, and the need for cognitive models, although no single solution to achieve widespread use of either was forthcoming.

Several areas require further research. Basic research is needed in planning and plan recognition, building natural-language interfaces, and applying architectures, such as blackboards to teaching systems.

From the viewpoint of communication, the symposium was a real success: discussion was lively and at times controversial. Research appears to be strong in depth, broad in perspective, and motivated by the promise of building more powerful teaching environments with greater knowledge, increased inference capability, and more complex reasoning ability. Researchers are very active and the field seems to be alive and well.

References

- Anderson, J.R. (1990) Analysis of Student Performance with the LISP tutor. In N. Frederickson et al. (Eds.), *Diagnostic Monitoring of Skill and Knowledge Acquisition*, Lawrence Erlbaum Associates, Hillsdale, NJ.
- Anderson, J.R., Boyle, C.F., Corbett, A.T. & Lewis, M.W. (1990) Cognitive modeling and intelligent tutoring. *Artificial Intelligence*, Vol. 42, 7-50.
- Bloom, B.S. (1984) The 2-Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-One Tutoring. *Educational Researcher*, Vol. 13., No. 6, pp. 4-16.
- Lesgold, A., Laijoie, S.P., Bunzo, M., and Eggan, G. (1990) A coached practice environment for an electronics troubleshooting job. In J. Larkin, R. Chabay and C. Sheric (Eds.), *Computer Assisted Instruction and Intelligent Tutoring Systems: Establishing Communication and Collaboration*, Lawrence Erlbaum Associates, Hillsdale, NJ.
- Lewis, M., McArthur, D., Stasz, C., & Zmuidzinas, M. (1990) Discovery-based Tutoring in Mathematics. In *Working Notes AAAI Artificial Intelligence Spring Symposium*, AAAI, Palo Alto, CA.

- Shute, V., Rose Garden Promises of Intelligent Tutoring Systems: Blossom or Thorn? (1990)
 Paper presented at the *Space Operations, Applications and Research (SOAR) Symposium*.
 Contact V. Shute, AFHRL, Brooks Air Force Base, TX 78235-5601.
- Shute, V.J., Glaser, R., and Raghavan, K. (1989) Inference and discovery in an exploratory laboratory. In P.L. Ackerman, R.J., Sternberg, and R. Glaser (Eds.), *Learning and Individual Differences*, W.H. Freeman, New York, pp. 279-326.
- Stevens, A., & Collins, A. (1977) The Goal Structure of a Socratic Tutor. In *Proceedings of the Association for Computing Machinery Annual Conference*. Also available as BBN Report No. 3518 from Bolt Beranek and Newman Inc., Cambridge, Mass., 02138
- Woolf, B. (1988) Intelligent Tutoring Systems: A Survey. In H. Shrobe and the American Association for Artificial Intelligence (Eds.), *Exploring Artificial Intelligence*. Morgan Kaufmann Publishers, Palo Alto, CA, pp. 1-43.